

BAYESIAN TOMOGRAPHY WITH PRIOR-KNOWLEDGE-BASED PARAMETRIZATION AND SURROGATE MODELING

G. A. Meles, N. Linde, S. Marelli



Data Sheet

Journal: Geophysical Journal International

Report Ref.: RSUQ-2022-003

Arxiv Ref.: <https://arxiv.org/abs/2201.02444> - [physics.geo-ph]

DOI: <https://doi.org/10.1093/gji/ggac214>

Date submitted: 23/12/2021

Date accepted: 31/05/2022

Bayesian tomography with prior-knowledge-based parametrization and surrogate modeling

Meles, G. A. and Linde, N. and Marelli, S.

June 13, 2022

Abstract

We present a Bayesian tomography framework operating with prior-knowledge-based parametrization that is accelerated by surrogate models. Standard high-fidelity forward solvers (e.g., Finite-Difference Time-Domain schemes) solve wave equations with natural spatial parametrizations based on fine discretization. Similar parametrizations, typically involving tens of thousand of variables, are usually employed to parameterize the subsurface in tomography applications. When the data do not allow to resolve details at such finely parameterized scales, it is often beneficial to instead rely on a prior-knowledge-based parametrization defined on a lower dimension domain (or manifold). Due to the increased identifiability in the reduced domain, the concomitant inversion is better constrained and generally faster.

We illustrate the potential of a prior-knowledge-based approach by considering ground penetrating radar (GPR) travel-time tomography in a crosshole configuration with synthetic data. An effective parametrization of the input (i.e., the permittivity distributions determining the slowness field) and compression in the output (i.e., the travel-time gathers) spaces are achieved via data-driven principal component decomposition based on random realizations of the prior Gaussian-process model with a truncation determined by the performances of the standard solver on the full and reduced model domains. To accelerate the inversion process, we employ a high-fidelity polynomial chaos expansion (PCE) surrogate model. We investigate the impact of the size of the training set on the performance of the PCE and show that a few hundreds design data sets is sufficient to provide reliable Markov chain Monte Carlo inversion at a fraction of the cost associated with a standard approach involving a fine discretization and physics-based forward solvers. Appropriate uncertainty quantification is achieved by reintroducing the truncated higher-order principle components in the original model space after inversion on the manifold and by adapting a likelihood function that accounts for the fact that the truncated higher-order components are not completely located in the null-space.

1 Introduction

Travel-time tomography refers to a large class of noninvasive imaging methods with applications comprising natural resource exploration (Taillandier et al., 2009), extraterrestrial investigations (Zhao et al., 2008; Khan et al., 2016), medical diagnosis (Stotzka et al., 2002) and non-destructive testing (Tant et al., 2018). In travel-time tomography, active or passive sources generate acoustic, elastic or electromagnetic waves, and the corresponding first-arrival times are used to infer properties of the illuminated medium (Rawlinson and Urvoy, 2006). High-resolution travel-time tomography of the shallow subsurface can be performed using ground-penetrating radar (GPR). GPR wave-propagation depends on the distribution of dielectric permittivity (ϵ) and electric conductivity (σ) in the subsurface. Permittivity and conductivity fields mainly determine wave propagation velocity and attenuation, respectively, while variations of either property result in scattering (Balanis, 2012). The penetration depth range varies greatly, ranging from few centimeters in wet bentonite clay up to several kilometers in ice, while in dry environments GPR waves can typically penetrate up to 30 meters (Slob et al., 2010). Under ideal conditions in terms

of coverage, signal-to-noise ratio and processing, GPR data can be used to obtain high-resolution estimates of permittivity and conductivity distributions, while simpler configurations and analyses allow for inexpensive identification/localization of subsurface discontinuities/anomalies (Grasmueck et al., 2005; Dou et al., 2016). Reflection surveys are relatively easy and inexpensive to deploy, but their performance is heavily affected by the limited penetration depth of GPR waves in conductive media (Piscitelli et al., 2007), such as soils. Cross-hole configurations are more demanding in terms of implementation as boreholes are needed, but are well suited for groundwater investigations (LaBrecque et al., 2002; Annan, 2005).

Deterministic GPR tomography can be carried out either by considering the travel-time and/or the amplitude (Olsson et al., 1992) of the first-arrival signals, or even considering the waveforms of the measured electromagnetic (EM) fields (Ernst et al., 2007; Kuroda et al., 2007; Meles et al., 2010, 2012), so-called full waveform inversion (FWI). Probabilistic inversion approaches based on Bayesian statistics have been proposed as an alternative, both for travel-time and FWI inversion (Nielsen et al., 2010; Cordua et al., 2012; Hunziker et al., 2019). A drawback of such search methods is the high computational cost associated with the many (often hundreds of thousands) solutions of the forward problem required to locate and sample from the posterior distribution; this can be computationally prohibitive even when state-of-the-art solvers are used (Warren et al., 2019). As an alternative to physics-based solvers, machine learning (ML) tools have become increasingly popular in recent years. A ridge regression ML algorithm has been introduced in GPR tomography to provide an optimal linear approximation for travel-times arrivals and several waveform attributes. Setting the problem as Gaussian and linear inevitably introduces modelling errors, but allows to describe the posterior analytically. Combining such a ridge regression linear model and proper treatment of modelling error is shown to provide better results than standard linear algorithms (Holm-Jensen and Hansen, 2020). Machine learning also offers a number of approaches to obtain inexpensive proxies mimicking non-linear input-output relations of standard forward solvers, and regressive neural networks have already been successfully used in the implementation of various global search GPR algorithms (Hansen and Cordua, 2017; Giannakis et al., 2019). In most inversion schemes, the input domain is parametrized relying on the same spatial discretization as the one used in standard forward solvers. When Finite-Difference Time-Domain (FDTD) schemes are employed, regular grids are usually considered and each grid point represents one inversion parameter. Standard 2D problems can then easily involve tens of thousands of parameters, leading to computationally extreme demands when relying on a Bayesian formulation of the inverse problem. Various strategies to reduce the effective dimensionality of the model domain have been attempted. For example, since the resolving power (in terms of wavelength) associated with the data is much lower than the discretization needed for accurate forward modelling, adjacent cells used for forward modeling can be clustered to represent a single inversion parameter. Such compression reduces the memory requirements of deterministic inversion schemes (Ernst et al., 2007) and decreases the ill-posedness of the corresponding problem, but parameterizations should ideally rather be based on available prior knowledge about the unknowns rather than modelling-specific constraints (Haario et al., 2004). Here we pursue an approach in which the effective dimensionality of the problem is determined by comparison of predictions of the the standard forward solver in the full model domain with respect to prior-knowledge-based parameters defining a lower-dimensional model domain (a manifold), such that the associated modeling errors are well below the expected noise level of the data.

To further decrease the computational costs associated with Markov chain Monte Carlo (MCMC) inversion, we employ unbiased surrogate modelling to approximate standard travel-time forward solvers (Xiu and Karniadakis, 2002). Various classes of surrogate models, such as those based on Gaussian process models or kriging (Sacks et al., 1989; Rasmussen, 2003) and polynomial chaos expansions (PCEs) (Xiu and Karniadakis, 2002; Blatman and Sudret, 2011), can be employed in Bayesian inverse problems (Nagel, 2019; Higdon et al., 2015; Marzouk and Xiu, 2009; Marzouk et al., 2007; Wagner et al., 2020, 2021a). In this contribution, we rely on

regression-based sparse PCE due to their extrapolation capabilities and robustness with respect to noise (Blatman and Sudret, 2011; Lüthen et al., 2021b; Marelli et al., 2021b). Several applications of PCE in the geosciences have been proposed. Examples include hydrogeological characterization of a three-dimensional multilayered aquifer, in which an efficient posterior exploration of a high-dimensional space is presented by Laloy et al. (2013), where the Karhunen-Loève transform is used to significantly reduce the dimensionality of the parameter space. A PCE model is then introduced to perform a two-stage MCMC inversion. PCE has also been used to automate velocity analysis in seismic processing, where the ideal velocity model can be defined as an optimizer of a variational integral in a semblance field. Calculating the relevant variational integrals is expensive and using probabilistic methods become computationally impractical. However, the variational integral can be replaced by a PCE, which results in massive reduction in overall computational costs (Abbasi and Gholami, 2017). PCE can also be used as a compression tool. In a magnetometric resistivity tomography study, the logarithm of the electrical conductivity field is compressed using few tens of Hermite polynomials rather than tens of thousands of discretization parameters, resulting in important computational gains (Vu et al., 2020).

In the context of Bayesian travel-time tomography, a pioneering use of PCE surrogates for 1.5D media has been recently introduced by Sochala et al. (2021). In their work, Sochala et al. (2021) demonstrate the potential of PCE trained using eikonal solutions in Bayesian travel-time tomography, assuming that the observation error dominates other sources of uncertainty. Here, we propose a Bayesian framework based on PCE to perform travel-time tomography of more complex 2D media. We show that PCE trained with FDTD data can be successfully designed and efficiently used for Bayesian inversion. Crucially, we also show that superior convergence properties are achieved when modelling error is taken into account in the inversion process.

2 Methodology

2.1 Bayesian inversion

Forward models are mathematical tools used to predict the outcome of physical experiments by operating on a set of input parameters. A forward problem expresses the relationship between input parameters and output values:

$$\mathcal{F}(\mathbf{u}) = \mathbf{y} + \epsilon, \quad (1)$$

with \mathcal{F} , \mathbf{u} , \mathbf{y} and ϵ standing for the physical law or forward operator, typically involving PDEs acting on locally defined quantities, the input parameters, the output and a noise term, respectively. The corresponding inverse problem aims at inferring properties of \mathbf{u} given data \mathbf{y} and any available prior information about \mathbf{u} . A general solution to an inverse problem is the posterior distribution defined over the input domain of the corresponding forward problem. This solution can be formally expressed using Bayes' theorem:

$$P(\mathbf{u}|\mathbf{y}) = \frac{P(\mathbf{y}|\mathbf{u})P(\mathbf{u})}{P(\mathbf{y})} = \frac{L(\mathbf{u})P(\mathbf{u})}{P(\mathbf{y})}, \quad (2)$$

where $P(\mathbf{u}|\mathbf{y})$ is the posterior distribution of the input parameter \mathbf{u} given the data \mathbf{y} , $P(\mathbf{y}|\mathbf{u})$ (also indicated as $\mathcal{L}(\mathbf{u})$) expresses the probability of observing the data \mathbf{y} given the input parameter \mathbf{u} , $P(\mathbf{u})$ the prior distribution in the input parameter domain and $P(\mathbf{y})$ the marginalized likelihood with respect to the input parameters (also known as evidence). While formally exact, Eq. (2) is seldom used when solving inverse problems as the computation of the evidence is in most cases very expensive. Instead, since the product $\mathcal{L}(\mathbf{u})P(\mathbf{u})$ is proportional to the posterior distribution and can be calculated for any value \mathbf{u} , one can use a Markov chain Monte Carlo (MCMC) method to draw samples from $P(\mathbf{u}|\mathbf{y})$ without considering the evidence. Here, we extract the most relevant features from the original input domain and perform the inversion on an effective

lower-dimensional manifold (see subsection 2.2). We also reduce the often excessive computational burden of high-dimensional MCMC problems by approximating the forward model via a much less demanding PCE surrogate (see subsection 2.3) while accounting for this approximation in the inversion.

2.2 Identification of an effective manifold

Without any loss of generality, the input in Eq. (1), can be formally expressed using new variables via a standard change of coordinates:

$$\mathcal{F}(\mathbf{u}(\mathbf{x}_{full})) = \mathbf{y} + \epsilon. \quad (3)$$

Such a reparameterization can help to shift the focus of the model representation from the requirements of the forward solver (local fine discretization) to one that is better suited for the Bayesian inverse framework in the sense that it encodes the prior distribution. The new coordinates \mathbf{x}_{full} can be defined, for example, in terms of Fourier coefficients, autoencoders, or principal components associated with random realizations of the generative model, which in turn can be related to the stochastic/topological properties inherent to the prior knowledge. Perhaps the most natural and widely used example is to discretize geostatistical prior knowledge on geostatistical hyper parameters into the Karhunen-Loève modes of a corresponding Gaussian random field. In this contribution, we rely instead on a data-driven approach using samples from a pre-existing generative model as this makes the method insensitive to any assumptions of the prior being a multi-variate Gaussian. In particular, we extract principal components (PCA) from a set of realizations of the prior model to identify the relevant non-local features of the input. Other transforms are possible provided it can be projected (transformed) on the forward-solver mesh.

In the case of PCA considered here, the forward problem can be exactly expressed as operating directly on the new set of coordinates, that is, the principal component coefficients:

$$\mathcal{M}(\mathbf{x}_{full}) = \mathbf{y} + \epsilon, \quad (4)$$

where $\mathcal{M} = \mathcal{F} \circ \mathbf{u}$, and \circ stands for function composition. Under the additional key assumption that the model can be faithfully represented by operating on an effective truncated M -dimensional subset \mathbf{x} of the new coordinates \mathbf{x}_{full} , we can write:

$$\mathcal{M}(\mathbf{x}) = \mathbf{y} + \hat{\epsilon}, \quad (5)$$

where $\hat{\epsilon}$ is a noise term including both observational noise and modeling errors related to the projection to the manifold. We assume that the first M components of \mathbf{x} include the most relevant aspects of \mathbf{x}_{full} , implying that we ignore the remaining coefficients by setting them to zero. However, we will include this error in the inversion by using an appropriate likelihood function that considers this error and other modeling errors induced by the surrogate modeling. Introducing a model operating on an effective set of coordinate \mathbf{x} casts the inverse problem on the M -dimensional manifold as:

$$P(\mathbf{x}|\mathbf{y}) = \frac{P(\mathbf{y}|\mathbf{x})P(\mathbf{x})}{P(\mathbf{y})}. \quad (6)$$

When \mathbf{x} is low-dimensional, the corresponding statistical prior distribution $P(\mathbf{x})$ can be *practically* estimated using random realizations of the generative model of $P(\mathbf{u})$. Many different MCMC sampling techniques could be used to draw samples from the posterior distribution $P(\mathbf{x}|\mathbf{y})$. Here, we use a Metropolis-Hastings algorithm based on a Gaussian proposal distribution with proposal scales determined by a diagonal covariance matrix. Using one common scaling factor, each element along the diagonal is proportional to the prior marginal variance associated with the

corresponding input parameter (Hastings, 1970; Marelli and Sudret, 2014). Since we assume that the impact of the ignored higher-order variables is non-zero, elements from this 'null space' defined by their prior distributions can be added after the inversion and used to properly assess uncertainty on \mathbf{u} .

A compressing procedure, consisting of a change of coordinates and a truncation, can also be made a-priori to reduce the dimensionality of the output domain. Note that the benefit of using compression in the output domain is associated with gains in storage and computational efficiency since each output variable requires its own PCE expansion (see section 3.1).

2.3 Polynomial chaos expansion

Forward models \mathcal{M} are typically implemented by potentially expensive schemes (e.g. finite element methods, FDTD, etc.) in order to accurately mimic the relationship between \mathbf{x} and \mathbf{y} in Eq. (5). A surrogate model $\tilde{\mathcal{M}}$ is a function that seeks to emulate the behaviour of an expensive forward model at negligible cost per run:

$$\tilde{\mathcal{M}}(\mathbf{x}) \approx \mathcal{M}(\mathbf{x}). \quad (7)$$

Among the many available classes of surrogate models, sparse adaptive polynomial chaos expansions are one of the most widely used due to their efficiency (the evaluation of the PCE polynomials can be done using sparse matrix multiplication), flexibility (e.g., in terms of polynomial definition and properties) and ease of deployment (Xiu and Karniadakis, 2002; Blatman and Sudret, 2011; Lüthen et al., 2021b; Métivier et al., 2020).

Polynomial chaos expansions are a type of stochastic spectral expansions that can be used to project the forward operator onto an orthonormal polynomial basis of a suitable functional space. Consider a stochastic Hilbert space \mathcal{H} equipped with the following scalar product:

$$\langle g(\mathbf{x}), h(\mathbf{x}) \rangle \equiv \mathbb{E}_{\mathbf{X}} [g \cdot h] = \int_{\Omega_{\mathbf{X}}} g(\mathbf{x})h(\mathbf{x})f_{\mathbf{X}}(\mathbf{x})d\mathbf{x}, \quad (8)$$

where the $\mathbb{E}_{\mathbf{X}}$ symbol denotes the statistical expectation operator with respect to the joint probabilistic distribution $f_{\mathbf{X}}(\mathbf{x})$ of the random vector of the input parameters \mathbf{X} . Then any forward operator of finite variance $\mathcal{M}(\mathbf{x}) \in \mathcal{H}$, that is,

$$\int_{\Omega_{\mathbf{X}}} \mathcal{M}(\mathbf{x})^2 f_{\mathbf{X}}(\mathbf{x})d\mathbf{x} \equiv \langle \mathcal{M}(\mathbf{x}), \mathcal{M}(\mathbf{x}) \rangle < +\infty, \quad (9)$$

can be expanded as a sum of polynomial basis elements Ψ_{α} of \mathcal{H} as (Xiu and Karniadakis, 2002):

$$\mathcal{M}(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}^M} a_{\alpha} \Psi_{\alpha}(\mathbf{x}), \quad (10)$$

where the multivariate polynomial basis functions Ψ_{α} are orthonormal with respect to the inner product in Eq. (8), that is:

$$\langle \Psi_{\alpha}, \Psi_{\beta} \rangle = \delta_{\alpha\beta}, \quad (11)$$

with $\delta_{\alpha\beta} = \prod_{i=1}^M \delta_{\alpha_i\beta_i}$. For practical purposes, the infinite series in Eq. (10) is truncated through a suitable truncation scheme, typically the maximum allowed polynomial degree p (Xiu and Karniadakis, 2002). For a review of more advanced basis truncation/construction schemes, the reader is redirected to Lüthen et al. (2021a). To calculate the set of expansion coefficients $a_{\alpha} = \langle \mathcal{M}(\mathbf{x}), \Psi_{\alpha} \rangle$, sparse regression techniques, also known as compressive sensing, have been demonstrated to be highly efficient in the context of surrogate modeling (Torre et al., 2019; Lüthen et al., 2021b,a), and are adopted in this paper. Once a_{α} is calculated from a finite data set, the

surrogate forward modelling predictor $\tilde{\mathcal{M}}(\mathbf{x})$ can be calculated at a negligible cost, by direct evaluation of Eq. (10). The accuracy of the surrogate model can be learned by comparing the performances of \mathcal{M} and $\tilde{\mathcal{M}}$ on a validation set if available, or by using cross-validation techniques (Blatman and Sudret, 2011; Marelli et al., 2021a).

2.4 Modeling-error aware Likelihood function

A PCE surrogate model can be used to faithfully predict physical quantities such as \mathbf{y} in Eq. (5). However, since we rely on an imperfect forward model operating on a truncated domain it is important to incorporate the corresponding error in the inversion process by using an appropriate likelihood function. The likelihood function discussed in Section 2.2 is assumed to be Gaussian, and can be designed to ignore or incorporate the modelling error associated with the PCE and the PCA truncation. We account for the modeling error by considering in the inversion process a covariance operator \mathbf{C}_D given by the sum of the covariance matrices describing data uncertainty, indicated in the following as \mathbf{C}_d , and a modelling term \mathbf{C}_{Tapp} . The covariance operator associated with the modelling error is defined in terms of a misfit matrix \mathbf{D}_e . For an unbiased model, \mathbf{D}_e consists of T realizations of the difference between the exact and the approximate forward responses. In case of biased models, the bias - indicated in the following as \mathbf{d}_d is added to the approximate responses (Hansen et al., 2014; Madsen and Hansen, 2018). The k^{th} column of \mathbf{D}_e is a n -dimensional vector corresponding to the single k^{th} realization of the modelling error, where n is the dimension of the output domain, thus, \mathbf{D}_e is a $n \times T$ matrix. Given the definition above, the covariance operator \mathbf{C}_D is defined as follows:

$$\mathbf{C}_D = \mathbf{C}_d + \mathbf{C}_{Tapp}, \quad (12)$$

where \mathbf{C}_{Tapp} is defined as the $n \times n$ matrix

$$\mathbf{C}_{Tapp} = \frac{1}{N} \mathbf{D}_e \mathbf{D}_e^T. \quad (13)$$

The likelihood incorporating the modelling error is then expressed as

$$\mathcal{L}(\mathbf{x}) = \left(\frac{1}{2\pi}\right)^{n/2} |\mathbf{C}_D|^{-1/2} \exp \left[-\frac{1}{2} (\tilde{\mathcal{M}}(\mathbf{x}) + \mathbf{d}_d - \mathbf{y})^T \mathbf{C}_D^{-1} (\tilde{\mathcal{M}}(\mathbf{x}) + \mathbf{d}_d - \mathbf{y}) \right]. \quad (14)$$

where $|\mathbf{C}_D|$ is the determinant of the covariance matrix \mathbf{C}_D . In practice, it is convenient that the training set itself is used to synthesise the matrix \mathbf{D}_e . The computation of the k^{th} column of the matrix \mathbf{D}_e is defined in terms of the output of the k^{th} element of the training set and a PCE approximation. This PCE estimate is based on cross-validation in which the results of the surrogate trained on the all the training input except the k^{th} one are used.

3 Application to GPR crosshole tomography

We apply now the proposed methodology to a GPR crosshole tomography problem. We consider for simplicity a GPR tomographic problem involving variations in permittivity, while assuming a constant and known electrical conductivity. We use the recording configuration displayed in Fig. 1(a) with 17 evenly spaced sources and receivers located in two vertically-oriented boreholes. The distance between the boreholes is 4.6 m, while the spacing between sources/receivers is 0.6 m. Each source is characterized by a 100 MHz Blackman–Harris pulse that is fired separately and the propagating signals are recorded at all receivers, such that a total of 289 travel-times are collected. We employ a 2D FDTD solver to simulate noise-free propagation in the transverse-electrical (TE) mode (Irving and Knight, 2006). Perfectly matched layers (PML) surrounding the propagation domain are used to prevent spurious reflections from contaminating the data, while appropriate

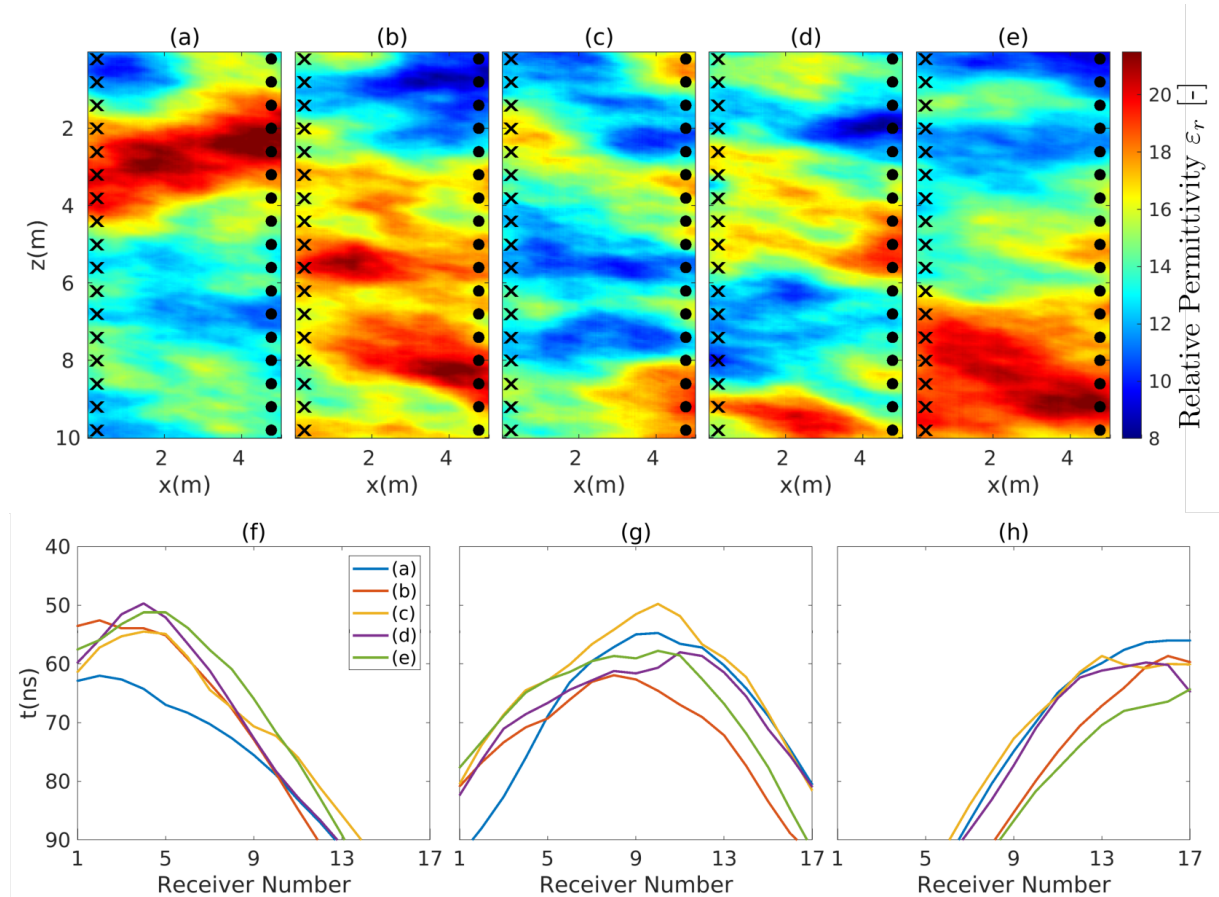


Figure 1: (a-e) Representative realizations of the generative model and (f-h) exemplary corresponding travel-time gathers. Crosses and circles stand for sources and receivers, respectively.

space-time grids are employed to avoid instability and dispersion artefacts. Travel-times are picked automatically based on a threshold associated with the relative maximum amplitude of each source-receiver pair. We rely on a generative model representing prior knowledge to create a set of relative permittivity distributions ϵ_r typical of a near-surface geophysical setting. We consider 5 m wide, 10 m deep random multi-Gaussian field realizations created using a pre-defined covariance structure based on the 2-D Matérn geostatistical model (Dietrich and Newsam, 1997; Laloy et al., 2015). The models are generated using six parameters: the mean (15), the standard deviation of ϵ_r (2.45), the anisotropy angle (85 degrees), the anisotropy ratio (0.3), the integral scale of the major axis (10 m) and a shape parameter (1.15). Figures 1(a-e) display representative samples of ϵ_r -distributions, together with a few representative gathers in Figs. 1(f-h). Given the permittivity range and the frequency content of the source wavelet, a spatial grid with $dx = dz = 4$ cm is used to avoid numerical dispersion and instability. Each generated permittivity field corresponds to a 125×250 dimensional space imposed by the FDTD grid.

3.1 Parametrization of input and output domains

Following a similar strategy to Laloy et al. (2013), we use the generative model to learn a prior-knowledge-based parametrization of the input domain. We create a total of 1000 random realizations from the prior (see Fig. 2). This number of realizations allowed inferring an appropriate sample mean and principal components of the underlying prior distribution, which led to convergence in terms of input domain root-mean-square error (in the following, RMSE) and structural similarity of the projections of distributions on subsets of principal components (Wang et al., 2004). Rather than considering the permittivity realizations to define the principal




Figure 2: (a-j) The first 10 principal components in the input domain. Crosses and circles stand for sources and receivers, respectively.

components, we take at each grid point the square root of each permittivity sample. This is done to simplify the PCE learning as the relationship between $\sqrt{\epsilon_r}$ and travel-times is linear in case of constant velocity. However for readability purposes comparison of projections and inversion results will be always referred to relative permittivity distributions. As expected, low wavenumber features are mainly associated with the first principal components, whereas small-scale details can be identified in higher modes (Fig. 2). In Fig. 3(a), a permittivity distribution is shown next to the approximate representations (Fig. 3(b-d)) obtained by projecting it on a moderate number of principal components. The remaining terms are ignored by setting their coefficients to 0, which corresponds to the mean value of their respective component. Table 1 summarizes the convergence in the relative permittivity domain of the various projections in terms of RMSE and structural similarity.

The similarity between Fig. 3(a) and Figs. 3(b-d) refer to proximity in the input domain, and does not guarantee that the map between input and output can be faithfully represented by operating on the reduced representations. The fidelity of such a reduced representation must, thus, also be assessed by considering the performance in the output domain. This aspect is highlighted in Figs. 3(e-g), showing representative travel-time gathers associated with the distributions in Fig. 3(a-d). Each approximate distribution seems to provide similar responses as those generated by the reference input. To quantify the agreement of the various reduced parametrizations, we consider 100 realizations of the generative model, and compute the resulting histograms of the travel-time residuals. The output domain root-mean-square error (in the following, *rmse*) of the

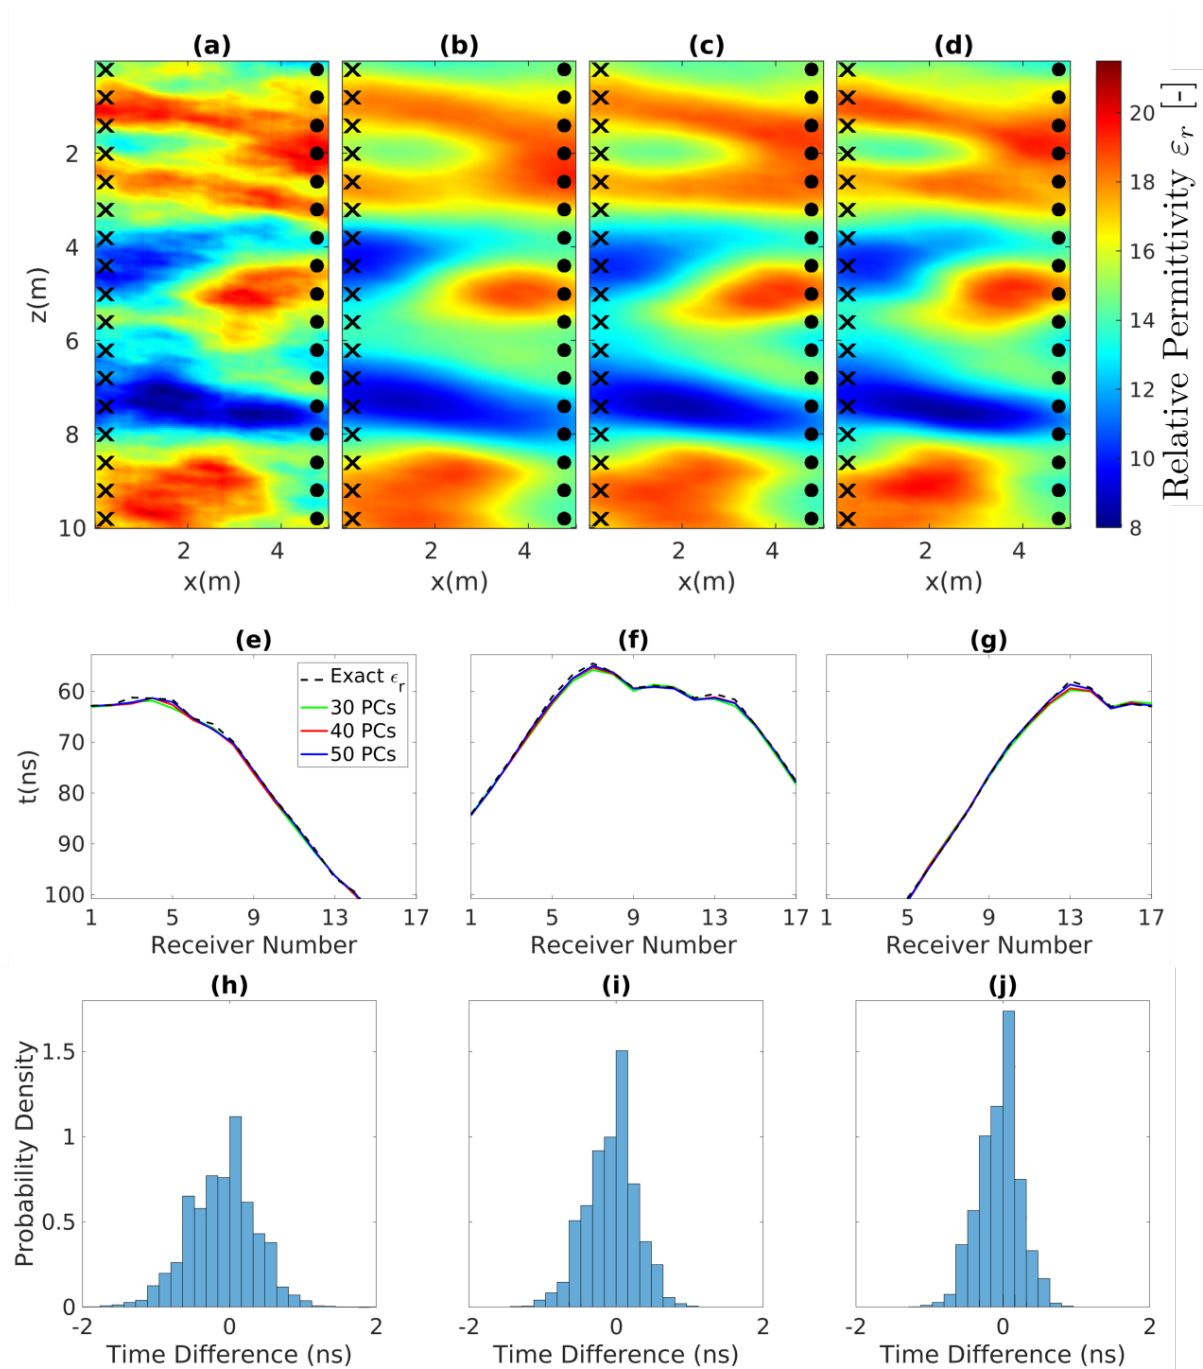


Figure 3: (a) The reference relative permittivity distribution used in the numerical inversion experiment and its representation using the first (b) 30, (c) 40 and (d) 50 principal components, and (e-g) exemplary corresponding travel-time gathers and (h-j) histograms of travel-time residuals based on FDTD forward simulations on the true field and reconstructions based on principal components. Crosses and circles stand for sources and receivers, respectively

Table 1: Assessment of the convergence of projections of permittivity realizations on a number of principal components in terms of input RMSE and SSIM. Note that the mean of the RMSE for a random realization of the generative model is ~ 15.2 .

PCA	RMSE ϵ_r	SSIM ϵ_r
30	0.88	0.70
40	0.80	0.72
50	0.69	0.75




Figure 4: (a) Sample mean and (b-d) representative principal components in the output domain

misfit between the data associated with the original distribution and its projections on 30, 40 and 50 principal components are 0.46, 0.34 and 0.28 ns, respectively, which are significantly smaller than the expected level of GPR data noise of 1 ns for 100 MHz data (Arcone et al., 1998). Based on this analysis, we can claim that the problem can be faithfully approximated involving forward simulations of fields with 30-50 of prior-knowledge-informed parameters.

Dimensionality reduction can also be exploited in the output domain to further reduce computational costs, storage and PCE training. We consider 1000 data sets corresponding to as many realizations of the generative model. We then compute via FDTD the 289 travel-times associated with the recording configuration discussed above. Based on these gathers we create a training set from which principal components are retrieved. The original data can then be projected on a subset of the principal components and characterized by the corresponding coefficients. The sample mean and representative principal components are shown in Fig. 4, while projections of selected source gathers are shown in Figs. 5(a-c). The agreement with the original data (black dashed line) and the projection on the first 10 (Fig. 5(a)) components transformed back into the original output domain is rather poor, but improves significantly when 20 (Fig. 5(b)) and 35 (Fig. 5(c)) components are used. Corresponding histograms of the travel-time residuals in Figs. 5(d-f) provide more quantitative insights concerning the performances of the compression. When only 10 principal components are employed the rmse is 0.94 ns. The fitting improves when 20 and 35 principal components are used, as the rmse reduces to 0.55 and 0.34 ns, respectively.

Despite the massive dimensionality reduction offered by PCA-decomposition, an FDTD-based Monte Carlo inversion would be extremely computationally demanding to carry out. To circumvent this and given the moderate number of input parameters, we replace the forward simulations with a PCE-based surrogate model, as implemented in the Matlab Package UQlab (Marelli and Sudret, 2014; Marelli et al., 2021a). Note that the implementation of a PCE requires a prior probability density function to build a corresponding set of orthonormal polynomials and

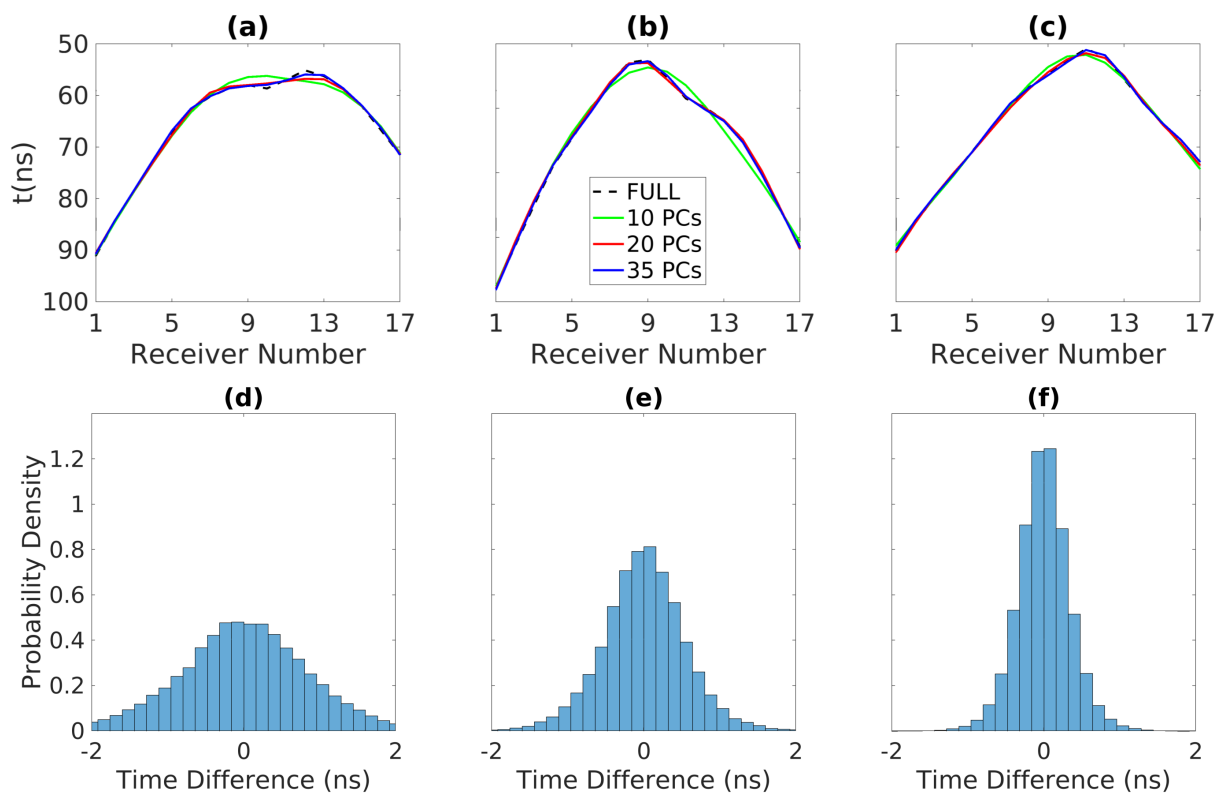


Figure 5: Decomposition in terms of (a) 10, (b) 20 and (c) 35 principal components in the Gather output domain of three representative gathers associated with a realization of the generative model, and (d-f) corresponding histograms of the residuals. The computation of the reference data is carried out using the FDTD solver.

the computation of a design set, consisting of travel-times associated with gathers computed using the FDTD scheme, to train the algorithm. We follow the approach described in [Torre et al. \(2019\)](#), by assuming statistical independence between the modes, and use kernel smoothing to estimate data-driven marginal PDFs. In particular, we re-use the realizations of the generative prior model (i.e., the 1000 realizations discussed above) to calculate the orthonormal polynomial basis.

Because PCE is designed for models with scalar outputs, the main benefit of compressing the multivariate forward model output is to reduce the number of PCEs that need to be trained and evaluated. The cost, in terms of computation of the training set, is the same with or without output compression, but the PCE training cost scales with the number of output variables. Using compression in the output then reduces the CPU time required to train the PCEs, while additionally reducing the overall memory fingerprint of the MCMC process.

The proposed implementation of PCE relies on peculiar representations of both input and output domains. To train the PCE via FDTD modelling we use a pixel-based parametrization of $\sqrt{\epsilon_r}$. In terms of PCE input variables, such a distribution is described via a set of 50 scalar components (i.e., the projection of each PC on a given $\sqrt{\epsilon_r}$ distribution), which offers little interpretability in terms of physical quantities. Therefore, unless stated differently, to offer easy to interpret images we present relative ϵ_r distributions in the pixel-based domains rather than the corresponding projections on principal components. This is done by transforming back a set of projections into the $\sqrt{\epsilon_r}$ domain, and then squaring that distribution at each point in space. Similar distinctions and transformation are important for output representation as well. FDTD modelling generates travel-time gathers which are projected on principal components before being used in the PCE process and inversion. The PCE output is defined in a domain that does not offer an immediate physical interpretation, but carry important statistical properties (see Appendix 1). To comment our results, in the following we will then use both the PCE output domain (in the following, 'PCE output') and the back-projection to the standard gather domain (in the following, 'Gather output').

3.2 PCE performance

Based on considerations of the expected noise level in field data and the performances of the proposed dimensionality reduction in the GPR problems discussed above, we perform PCE using 50 principal components in the input and 35 in the PCE output domain. With these settings, we analyze the PCE performance as a function of the training set size. In this study, the output is trained to predict the projection coefficients of travel-times corresponding to all source-receiver pairs using a size of \mathbf{x} (M in Eq. (10)) that is 50. A final polynomial degree, p , of 5 was automatically selected by the adaptive sparse expansion algorithm from [Blatman and Sudret \(2011\)](#); [Lüthen et al. \(2021b\)](#). Histograms of travel-time residuals between the reference travel-time data and PCE results obtained using 50 principal components in the input domain and 50 (Fig. 6(a)), 500 (Fig. 6(b)) and 900 (Fig. 6(c)) training sets. When only 50 training sets are used, the rmse in the Gather output is 0.86 ns (Fig. 6(a)). With an increasing size of the training set, the PCE performances improve significantly, with rmse of the misfit in the Gather output reducing to 0.51 and 0.47 when 500 (Fig. 6(b)) and 900 (Fig. 6(c)) data sets are employed, respectively. Note that regardless of the number of training data sets, the PCE results are unbiased and quite closely resemble Gaussian distributions.

By using only 50 principal components to characterize the input domain for the PCE predictions, we assume that higher-order principal components belong to the null space of the forward problem. While not strictly valid, this assumption is partially justified by the histograms in Figs. 3(j) and 5(f).

For comparison purposes, we now show the corresponding histogram of the travel-time residuals obtained when using the straight-line-approximation (Fig. 6(d)). As expected, the straight-line-approximation always overestimates the exact travel-times, and has an rmse of 1.53




Figure 6: (a) Histograms of the model errors with respect to the PCE prediction when using 50 and 35 principal components in the input and output domain, respectively, and 50 training sets for the expansion. (b) and (c) are constructed with the same parameters but with 500 and 900 training sets, respectively. (d) Histograms of the model errors associated with the straight-line-approximation. The computation of the reference data is carried out using the FDTD solver.

ns. By removing the mean value of the residuals (bias), the rmse reduces to 0.52 ns, which is still higher than for the PCE results using a training set size of 500 and above. In contrast to the PCE results, the straight-line-approximation histogram is clearly asymmetrical and, hence, non-Gaussian, which makes the likelihood in Eq. (14) inaccurate as it is based on a Gaussian assumption. The computational burden involved in constructing the PCE is comparable to what is required to account for modelling error for the straight-line-approximation (i.e., the evaluation of FDTD training sets) and its evaluation cost is negligible with respect to the computation time of the reference FDTD solver.

3.3 Inversion results

In section 3.2, we have shown how a PCE can approximate a forward FDTD solver in the computation of travel-times in a realistic geophysical scenario operating on a reduced input domain. Importantly, its performance using 500 training sets lead to modeling errors that are much less than the expected data noise and it is notably better than straight-line approximations that are biased, have skewed error statistics with larger variance. Here, we explore the potential of PCEs based on 50 PCs in the input domain (obtained using 50, 500 and 900 training sets, respectively) and a straight-line solver in performing probabilistic MCMC inversion. The entire inversion process was carried out using the UQLab Matlab package framework (Wagner et al., 2021b). We invert for the permittivity distribution shown in Fig. 3(a), which is used without model compression to generate travel-times using the reference FDTD solver. Uncorrelated Gaussian noise characterized by $\sigma^2 = 1ns^2$ is added to the FDTD data. Compression in the output domain is then applied, and the resulting data are inverted using the PCE surrogate model. We use a Metropolis-Hastings algorithm, and run 15 Markov chains in parallel for 1.5×10^5 iterations. Based on the visual inspection of the chains, we discard the first 1000 iterations as burn in for the evaluation of mean and variance values of the final results. The scaling factor of the proposal distribution is tuned so that an acceptance rate of about 30% is achieved for each experiment. Inversion results in terms of mean and standard deviations ignoring the model error (i.e. setting both $\mathbf{C}_{T_{app}} = 0$ and $\mathbf{d}_d = 0$ in Eqs. (12) and (14), respectively) are shown in Fig. 7. The PCE and Gather output rmse associated with the PCE-simulations of the posterior mean results in Figs. 7(a-c) are 2.5 and 1.44 ns for a training set of 50, 1.24 and 1.18 ns for a training set of 500 and 1.10 and 1.19 ns for a training set of 900, respectively. The slightly worse result associated with the Gather output for the PCE trained with 900 data sets is strictly speaking not indicative of a worst inversion, since the cost function is defined in terms of PCE output data. The rmse associated with the posterior mean obtained with the straight-line-approximation (Fig.

7(d)) are significantly higher at 4.70 and 2.05ns in PCE and Gather output, respectively.

Figures 7(e-g) show the corresponding standard deviations of the posterior distributions without considering the variability in the truncated PCA coefficients (the assumed null-space). High standard deviations are mainly located along the boreholes and the top and the bottom of the model due to the limited sensitivity of the data with respect to those areas. However, an accurate comparison of posterior uncertainty in the input domain can only be carried out if we consider elements from the null space. The corresponding standard deviation distributions obtained when adding the truncated PCA components (the assumed null-space)(Fig. 7(i-l)) indicate an overall increase in posterior standard deviation ($\sim +20\%$), thus, demonstrating the importance of adding elements from the 'null space' to properly assess posterior uncertainty.

We now run another set of inversions taking into account the modelling error to appreciate its effect on the corresponding solution. Fig. 8 shows covariance matrices \mathbf{C}_{Tapp} of the modeling errors for the PCE and straight-line solver used in the previous example. Note that in all the covariance matrices presented in Fig. 8 the off-diagonal elements are significantly smaller than the diagonal ones, thus indicating poorly correlated modelling error. This feature, common to both the PCE and the straight-line models, is due to the output compression. The corresponding PCE and Gather output rmse values of the PCE- simulated data are 2.25 and 1.41 ns (Fig. 9(a)), 1.26 and 1.16 ns (Fig. 9(b)), 1.07 and 1.17 ns (Fig. 9(c)), rather similar to those of the previous example. However, the rmse for the straight-line inversion are significantly smaller at 1.54 and 1.26 ns, for PCE and Gather output, respectively, (Fig. 9(d)), mainly as a consequence of the bias removal in the likelihood function. Standard deviation results are summarized in Fig. 9(e-l) and show a similar trend to what discussed about Fig. 7(e-l).

The results discussed above show that differences between PCE and straight-line-approximation-based inversion can be appreciated in terms of data misfit. Rather interestingly, ignoring or incorporating modelling error in PCE inversion seems to have little impact on the results in the input domain. The reason is most likely that the PCE is bias free and that the model error is significantly smaller than the observational errors. The PCE-based posterior mean fields in Figs. 7(a-c) agree with the projection of the reference distribution on the first 50 principal components (Fig. 3(d)) similarly to what the results in Figs. 9(a-c)) do. However, the two sets of solutions differ in terms of MCMC convergence, as assessed using the potential-scale reduction factor \hat{R} considering the variance of the individual Markov chains with the variance of all the Markov chains merged together (Gelman and Rubin, 1992). Convergence is usually assumed if $R < 1.2$ for all parameters. In our experiments, when modelling error is taken into account, all of the 50 parameters have a value of R smaller than 1.2 for any choice of the training set dimension size. However, if modelling error is not taken into account, R is larger than 1.2 for at least some of the parameters, regardless of the size of the training set. By contrast to the PCE results, the impact of incorporating model error in the inversion is very relevant when the (biased) straight-line solver is employed, as significantly better results are achieved when \mathbf{C}_{Tapp} and \mathbf{d}_a are taken into account (compare the poor agreement of 7(d) and the good agreement of Fig. 9 with respect to Fig. 3(d)). As for the previous cases, convergence is achieved only if model error is taken into account.

A more quantitative assessment can be achieved by comparing the reference input and the corresponding inversion solutions in terms of RMSE, structural similarity and Scoring values in the input domain (Gneiting and Raftery, 2007; Levy et al., 2021). We consider the posterior mean from the posterior distribution, including elements from the null space, and compute the RMSE and structural similarity for inversion based on PCE and straight-line-approximation trained on 900 sets with and without modelling error. To further assess both the statistical consistency between predictions and observations (calibration) and the sharpness of the prediction, we use the MCMC samples, transformed back in the pixel domain, to obtain a Gaussian approximation \hat{P} of the estimated posterior PDF and use it to compute the logarithmic score with respect to the true value \mathbf{u}_{true} at each point, that is, $\log S(\hat{P}, \mathbf{u}_{true}) = -\log(\hat{P}(\mathbf{u}_{true}))$. Low values of the logarithmic score are associated with PDFs under which the true value has high probability




Figure 7: MCMC-based posterior mean obtained while ignoring modeling errors in the likelihood function with PCE based on (a) 50, (b) 500 and (c) 900 training sets and (d) using a straight-line-approximation solver. The PCE operate between reduced domains defined in terms of 50 and 35 principal components associated with $\sqrt{\epsilon_r}$ and travel-time gathers, respectively. Results are shown as ϵ_r distribution for readability purposes. Corresponding posterior standard deviations (e-h) without considering the null-space and (i-l) considering the null-space. Crosses and circles stand for sources and receivers, respectively




Figure 8: Model error covariance matrices associated with PCE based on (a) 50, (b) 500 and (c) 900 training sets. (d) Model error covariance matrix associated with the straight-line solver, based on 900 training sets

Table 2: Assessment of the inversion results in the input and output domains for PCE and straight-line-approximation-based inversion, with or without modelling error included in the definition of the likelihood function.

Model	RMSE Mean ϵ_r	SSIM Mean ϵ_r	Mean LogS ϵ_r	rmse Mean PCE	rmse Mean Gather
PCE	1.10	0.69	1.63	1.07 ns	1.17 ns
PCE NO	1.09	0.69	1.63	1.10 ns	1.19 ns
SL	1.23	0.66	2.13	1.54 ns	1.26 ns
SL NO	1.56	0.63	3.72	4.70 ns	2.05 ns

(Good, 1992; Friedli et al., 2021). Scoring values for inversion results associated with PCE and straight-line-approximation solvers including and ignoring model error are shown in Fig. 10. The results of the fitting performances of the inversion results in the input and output domains are summarized in Table 2. PCE-based inversions, with or without modelling error (PCE and PCE NO in Table 2, respectively) undeniably perform better than the corresponding schemes using the straight-line-approximation (SL and SL NO, respectively). The difference is substantial when RMSE in the input domain is considered, while the discrepancy becomes less important when structural similarity is used as a measure. Note that, as also observed by Levy et al. (2021), including the modelling error in the inversion has as expected a minor adverse effect on fitting in the input domain regardless of the modelling scheme used as a consequence of the more inflated likelihood function. Performances differ very clearly in favor of PCE when scoring values are considered. Both PCE results (Figs. 10(a-b)) provide significantly smaller values of the logarithmic score in large portions of the domain with respect to the straight-line solutions (Figs. 10(c-d)). The differences in-between the PCE results are very minor, indicating only an extremely small improvement when modelling error is taken into account (see yellow arrows in Figs. 10(b)). By contrast, including modelling errors has a dramatic and beneficial impact on the straight-line solutions (see white arrows in Fig. 10(d) and compare with Fig. 10(c)).

From here onward, we exclusively consider the PCE-results associated with 900 training sets including modelling error (see Fig. 9(c)) as our reference inversion result. For these results, we show posterior realizations based on the reduced model parameterizations (Figs. 11a-e). These images are very smoothly varying property fields and they do not express the posterior variability of the original problem. In contrast, Figs. 11f-j show samples from the full posterior distribution obtained by including elements from the assumed null space (that is from the prior distribution of the higher-order truncated PCA coefficients). We see that these latter realizations bring in small-scale variability that can be highly important in subsurface applications particularly with




Figure 9: MCMC-based posterior mean obtained while incorporating modeling errors in the likelihood function with PCE based on (a) 50, (b) 500 and (c) 900 training sets and (d) using a straight-line-approximation solver, corresponding posterior standard deviations (e-h) without considering the null-space and (i-l) considering the null-space. Crosses and circles stand for sources and receivers, respectively

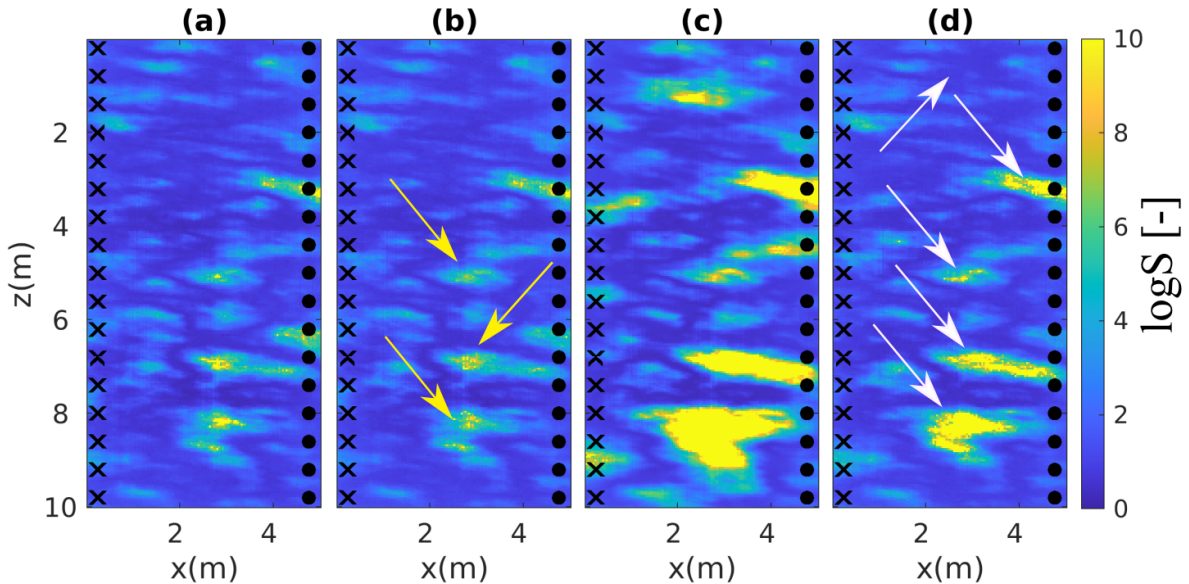


Figure 10: Logarithmic score values with PCE (a) ignoring model error and (b) considering the model error in the likelihood function. Corresponding values for the straight-line-approximation (c) ignoring and (d) considering the model error in the likelihood function. Crosses and circles stand for sources and receivers, respectively

respect to flow- and transport applications, but that are unresolvable by the data.

The quality of the inversion results can be further analysed by looking at PCE evaluations on ensembles consisting of a number of random realizations from the prior and the posterior distributions. The resulting distributions obtained by applying the PCE on the prior and posterior samples (known as prior and posterior predictive distributions) allows us to appreciate the reduction in uncertainty in the output domain of the posterior with respect to the prior. In this case we consider 10000 realizations both of the prior and posterior, and summarize the corresponding output values in Fig. 12 along the observed data and the PCE evaluated at the posterior mean. Not only is the match between the predictions based on the posterior mean and the observed data good, but we can also appreciate the much thinner posterior predictive distributions compared to the prior predictive distributions. This analysis in the PCA output domain does not offer an immediate physical interpretation, but we can transform these predictions back into the more natural Gather output domain. This is done in Fig. 14 showing, for each source/receiver pair considered the much smaller width of the posterior compared with the prior predictive distributions. Finally, we explore the inversion results in the PCA-based input domain by considering scatter plots for a number of representative input principal components (see Fig. 13). The data set is again shown to be highly informative in reducing posterior with respect to prior uncertainty. Nevertheless, we can also see that some of the inverted coefficients, for example, the 8th (see Fig. 2(h)), are not well resolved. This is not a consequence of the PCE, but of the intrinsic limitations of borehole configurations, allowing for compensation between high and low slowness values in explaining travel-times and, therefore, resulting in low sensitivity with respect to certain principal components. Compared to a deterministic inversion procedure, these limitations are expressed in terms of quantifiable uncertainty and can be interrogated as such.

4 Discussion

A new parametrization of the inverse problem can be attractive, even if forward solvers are applied on classical grids. One previous example of such an approach for GPR inversion in terms




Figure 11: Samples from the posterior distribution (a-e) ignoring the null-space and (f-j) corresponding realizations including random realizations from the null space. Crosses and circles stand for sources and receivers, respectively

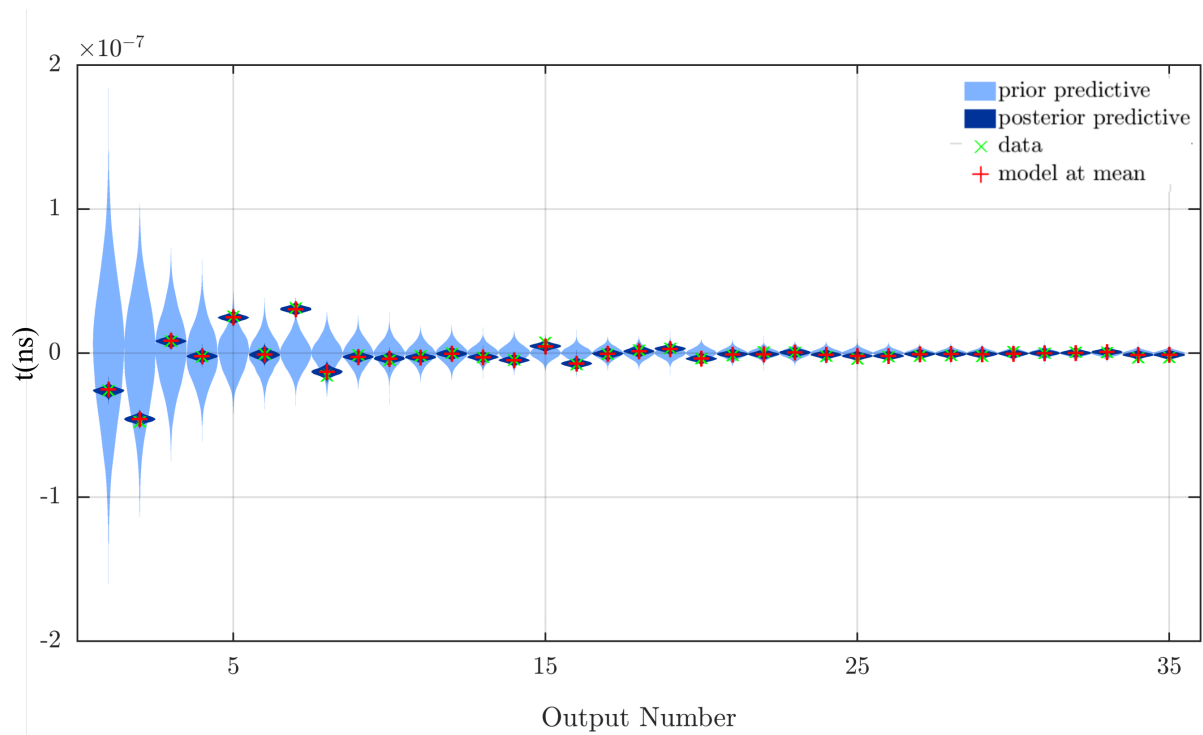


Figure 12: Prior and posterior predictive distributions in the PCE output domain. Light and dark blue histograms in the output domain correspond to 10000 realization of prior and posterior distributions, respectively. The green 'x' markers indicate the 'observed' data associated with the reference model, whereas the red '+' indicate the application of the PCE to the mean of the posterior.

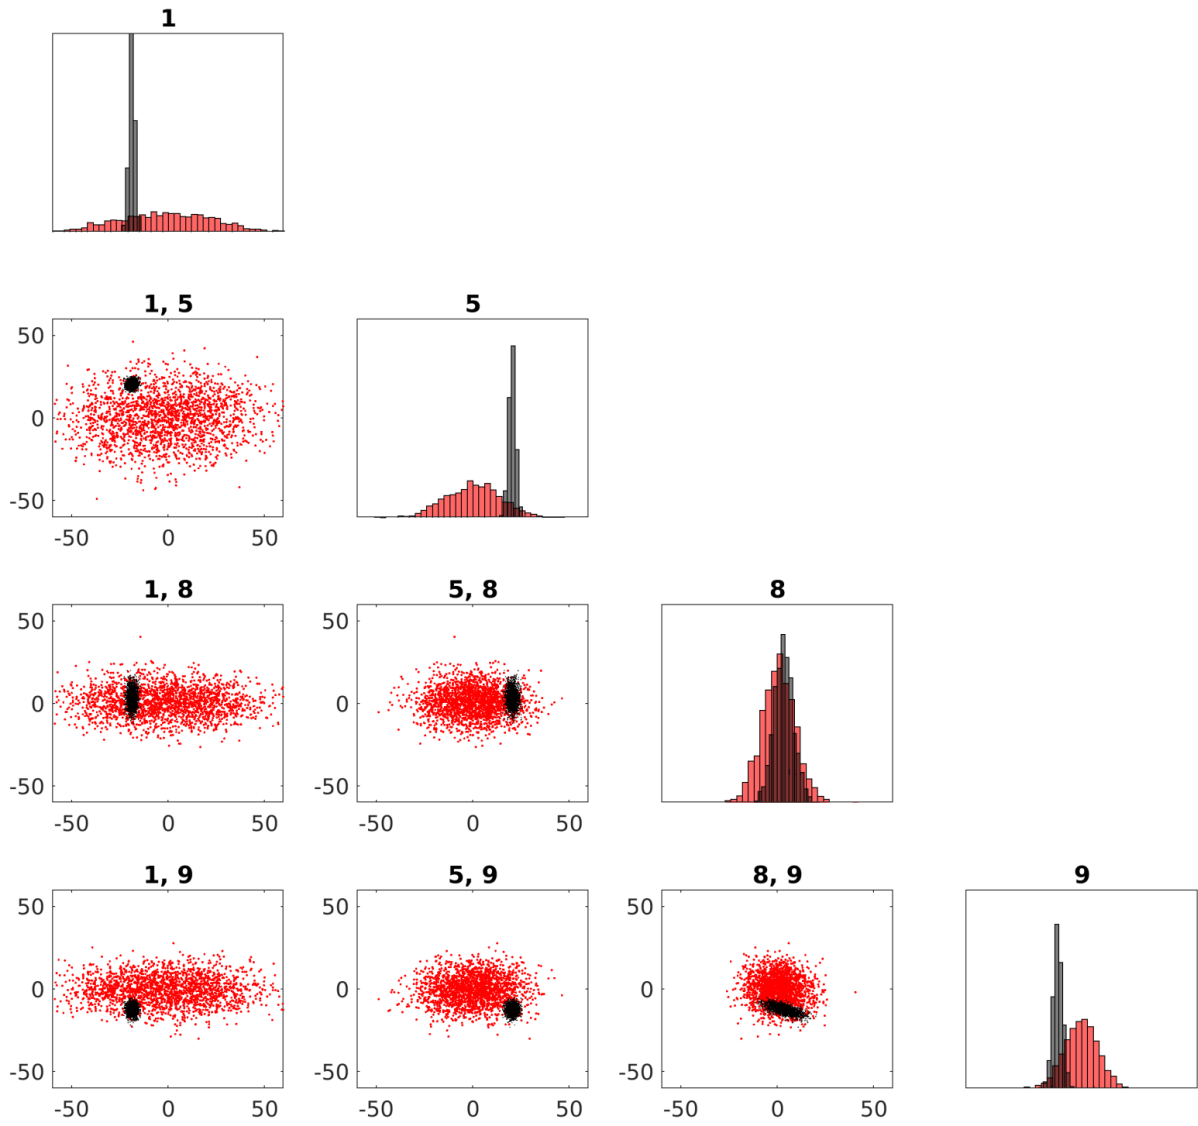


Figure 13: Scatter plot representation of prior (red) and posterior (black) distributions for a set of input variables.

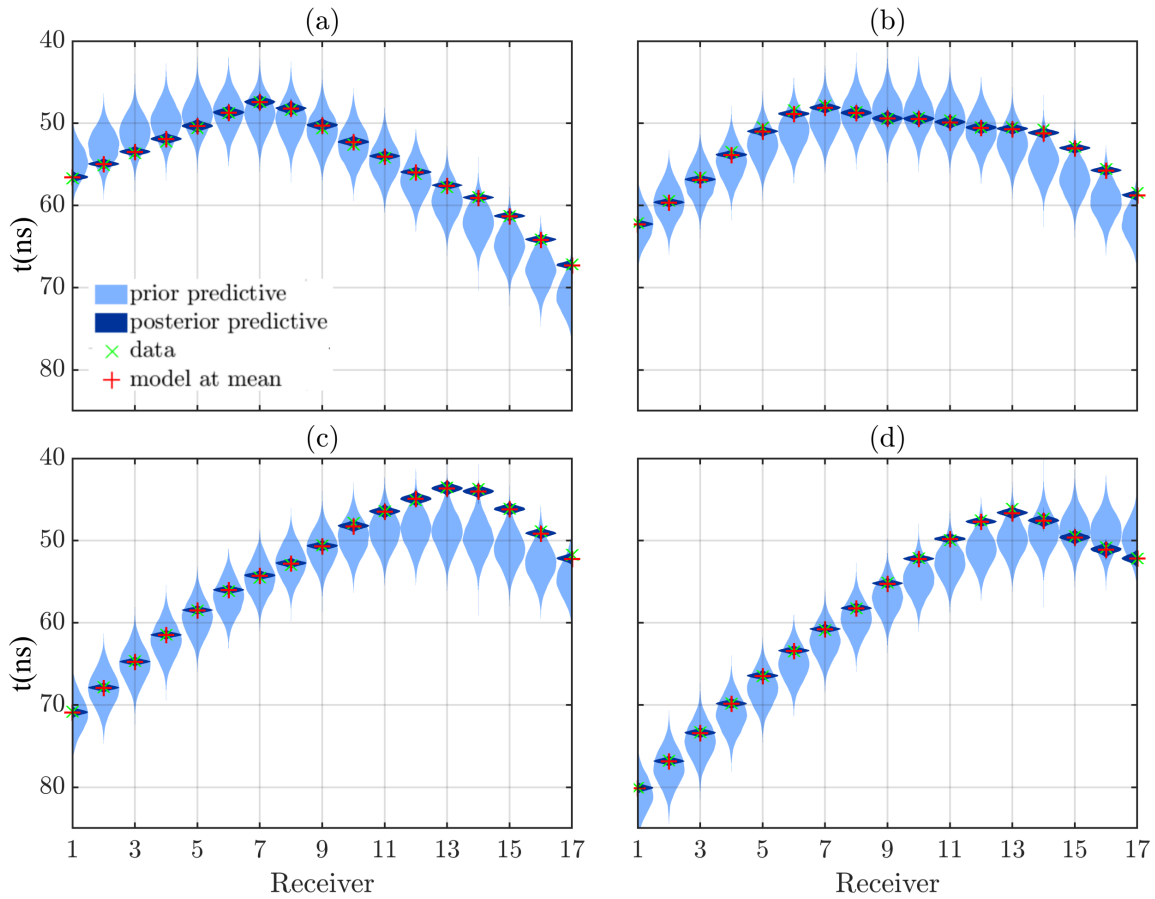


Figure 14: (a-d) Exemplary prior and posterior predictive distributions in the travel-time Gathers output domain. Differently than the PCE output domain (see Fig. 12), these diagrams offer an immediate physical interpretation. We can appreciate, for example, that the width of the posterior predictive distributions varies as a function of the travel-time.

of hydrogeological applications is [Giannakis et al. \(2021\)](#). In their work, they use randomly generated anisotropic 2D fractals to populate a training set, whose principal components are used to approximate fractal and non-fractal soil distributions. A heuristic FDTD algorithm, subsequently minimizing a full-waveform cost function along a few tens of ordered principal components, was shown to provide reliable and relatively inexpensive results by operating via FDTD on a massively reduced inversion domain. In our case, despite the strong decrease in complexity offered by a PCA-based parametrization, using the FDTD solver on the reduced input domain would involve extremely demanding computational costs and MCMC inversion would still be unpractical. We have circumvented this limitation by coupling an appropriate prior-informed parameterisation with surrogate modelling.

Surrogate modelling has emerged as a computational paradigm over the last 20 years. The massive decrease of the computational effort associated with such modelling is achieved by approximating the underlying computational model with a simple and easy-to-evaluate function. A key requirement to allow implementation of PCE is that the number of input variables describing the computational model is relatively small (i.e., up to a few tens). As the number of input variables increases, the number of coefficients required to define the surrogate model grows polynomially in both the size of the input and the maximum degree of the expansion, leading to non tractable problems if the model response is highly nonlinear in its input parameters ([Torre et al., 2019](#)). When a generative Gaussian random field is considered, PCs are uncorrelated and normally distributed. When considering general non-Gaussian random fields as priors, they can still be represented by their PCs, at the cost of no longer being Gaussian and independent. Nevertheless, by construction the PCs will always have uncorrelated marginal distributions.

In our approach we allow for arbitrary marginal distributions, hence extending its usability far beyond Gaussian priors, but independence is enforced, rather than just correlation. This is a limitation that can be noticed when the prior is a constrained random field, for which non-trivial trade-offs between PCs can then be observed. In such cases, alternative low-dimensional representations can be used, either by modifying assumptions on the dependence structure of the reduced space components (e.g., using custom copulas or multidimensional kernel smoothing) or by changing the “dimensionality reduction” step. Generally speaking, dimensionality reduction could be achieved by other means than PCA. Other parametrizations (wavelets, Fourier coefficients, Slepian functions, spherical harmonics, latent variables in deep generative models), tailored for specific applications, may be employed to express the effective dimensionality of the input domain. One setting in which PCE could be particularly effective, is in studies where only few key parameters (such as position, shape and basic physical properties of scatterers or layers) properly characterize the target medium.

Once a parametrization has been found to reduce the effective dimensionality of the input domain and a high fidelity PCE has been designed, inexpensive and effective MCMC inversion can be implemented. We have also shown how to include the impact of the ignored variables approximating the ‘null space’ in the likelihood function and in expressing the full posterior uncertainty.

We have relied on a standard Metropolis-Hastings sampling algorithm, but alternative more effective sampling methods, such as Adaptive Metropolis, Hamiltonian Monte Carlo, or the Affine Invariant Ensemble Algorithm ([Duane et al., 1987](#); [Haario et al., 2001](#); [Goodman and Weare, 2010](#)) could be easily implemented in our workflow.

When considering modelling error in the likelihood function, our results always provided converging solutions with respect to Gelman-Rubin diagnostics, both for PCE- and straight-line-based inversion. Neglecting modeling error had a highly negative impact on inversions based on the (biased) straight-line approach in terms of RMSE, SSIM and LogS. The results obtained when including modeling errors in the likelihood function still provided less good results than the PCE-based forward modeling, most likely due to a larger variance and a skewed distributions implying that the Gaussian assumption in the likelihood is questionable. The effect of including

the model error in the PCE-based inversion was very small, probably due to its unbiased nature, the relatively small variance and the Gaussian-like and symmetric error statistics.

Most of the computational costs associated with MCMC inversion based on PCE is due to the training of the PCE surrogate, which in our case required the computation of several hundreds of FDTD data sets. However, the re-training of a PCE model is only needed when a different combination of forward model, prior model or experimental set-up are required. Also, it should be noted that these training sets of up to 1000 data sets is very small compared to what is used in deep-learning-based proxy modeling. This makes this surrogate-modeling-based approach particularly attractive, for example, for monitoring applications where both solver and prior information remain unchanged between experimental campaigns.

Even if this was not necessary for the tests presented herein, re-training of a PCE may be needed during an MCMC inversion when considering more non-linear problems. In travel-time computation, the degree of non-linearity of the forward problem depends largely on the variance of the slowness field. For the class of permittivity distributions we considered here, the straight-line-approximation is poor (see Fig. 9(d)) indicating a certain degree of non-linearity. Even so, a single surrogate model provided good results on input samples realistically representing the entire class of permittivity distributions involved (see Fig. 6c). Therefore, a single PCE was used during the entire inversion process. However, when the forward model is more non-linear, constructing a globally-valid surrogate applicable for any realization of the prior can be challenging. In these cases, the accuracy of the surrogate can be refined in regions of high posterior probability (Li and Marzouk, 2014; Wagner et al., 2021a). Even if we considered borehole GPR applications herein, the the strategy could be easily adapted to other imaging problems such as active or passive seismic tomography at different scales (Bodin and Sambridge, 2009; Galetti et al., 2017).

In this preliminary study, two important assumptions were made. First of all, we assumed electrical conductivity to be constant throughout the inversion domain. In theory, PCEs accounting for the effects of varying conductivity distributions on travel-times (or waves amplitude), could be designed. However, the number of input/output variables would increase, and so the complexity of the entire process describe here. Secondly, we only considered synthetic data contaminated with uncorrelated Gaussian noise. Application of the proposed workflow to field data will face additional difficulties associated with the inclusion of actual noise in the inversion process. While the design of the surrogate only relies on synthetic data, simulation of an appropriate training set is challenging, as it would require 3D FDTD simulation and proper modelling of the source mechanism (e.g., for the GPR case, adequate antenna simulation (Warren et al. (2016))). Note that for field application, a proper treatment of conductivity would be necessary.

Finally, a note on computational cost. All the numerical experiments discussed in this paper have been executed on a workstation equipped with 16GB DDR4 of RAM and powered by a 3.5GHz Quad-Core processor running Matlab on Linux. With this configuration, the 15 chains / 1.5×10^5 iterations MCMC simulations presented here took around 6 hours to be executed. These performances indicate that PCE has the potential to allow MCMC inversion to be carried out without relying on massive computational resources.

5 Conclusions

Model parameterizations for MCMC-based inversions are ideally informed by prior knowledge, as this allows for massive dimensionality reductions compared with parameterization based on the grids used in standard forward simulations. In a crosshole GPR setting, we have demonstrated how travel-time modelling using 30-50 prior-knowledge-based parameters based on PCA decomposition, instead of more than 30'000 parameters if relying on the FDTD grid, can predict the output well below the expected observational noise-level. We also demonstrate that high-quality PCA-based decomposition can be achieved in the output domain (from 289 travel-times to some tenths of PCA-coefficients). This massive reduction of the effective size of the input and output domain

allows for efficient implementation of a PCE surrogate model using less than 1000 training sets. The resulting PCE surrogate model outperforms a straight-line-based solver that produces biased, higher variance and skewed model predictions. More importantly, the associated PCE-based modeling errors with respect to a FDTD solver are much smaller than the typical observational errors. Thus, allowing efficient replacement of FDTD solvers in the context of crosshole GPR travel-time tomography. This model reduction in the input and output domains combined with PCE-based modeling have dramatically reduced the computational burden of our problem, allowing us to successfully perform non-linear MCMC inversion at a low computational cost. Even if the PCE-based modeling error is low, we show how to account for it in the likelihood function and demonstrate that this inclusion is needed to obtain converging MCMC chains. We further highlight that the truncated higher-order PCA components that are ignored in the dimension-reduced inversion need to be re-introduced when assessing posterior uncertainty and for obtaining realistic-looking posterior realizations. Indeed, even if the small-scale features cannot be resolved by the data, thereby, allowing them to be left out of the inversion process, they must be reintroduced to solve the original inverse problem (not letting the posterior distribution be affected by the dimensionality reduction) and also as small-scale unresolvable features might be important when using the results for predictions in terms of, for instance, hydrogeophysical applications involving transport processes that are highly impacted by small-scale heterogeneity. We stress that these results are expected to be easily extendable to other types of wave-based physics (e.g., seismics) also at larger-scales or in 3D, while applications to diffusion-based geophysics (e.g., electrical resistance tomography) would need to be investigated in detail in future studies.

6 Acknowledgments

We thank the editor, Frederik Simons, and the reviewers, Jan Dettmer and Thomas Mejer Hansen, for their insightful suggestions and comments that have helped us improve the clarity and quality of this manuscript.

7 Data availability

The data underlying this article will be shared on reasonable request to the corresponding author.

Appendix A

Compression in the output domain is convenient when considering PCE-based surrogate modeling, as every output requires its own expansion which, in the MCMC sampling, is computed independently from the others. A reduction in output dimension imply, therefore, a decrease in computational costs and memory requirements, which can be particularly important when very long MCMC chains are used. However, switching to a different coordinate set poses the problem of determining the statistical nature of the new output variables. We demonstrate that the complete and reduced output variables share the same relevant properties; this is a consequence of the linearity of PCA. In the following, N and n represent the dimension of the complete and reduced output domains (in our considered example, $N=289$, $n=35$). The standard output domain of travel-times, indicated here as $W_i, i \in [1, N]$, are assumed to constitute a set of independent identically-distributed Gaussian variables characterized by $\sigma^2 = 1ns^2$. The new set $Y_j, j \in [1, n]$ of output variables consists by definition of linear combinations of the random variables W_i , that is:

$$Y_j = a_{ij}(W_i - \bar{W}_i), j \in [1, n], \quad (15)$$

where \bar{W}_i and a_{ij} stand for the sample mean and the elements of the orthogonal matrix representing the principal components, respectively (see Fig. 4), and summation over repeated indices is

implied. Each $Y_j, j \in [1, n]$ is given by the sum of Gaussian random variables and is, therefore, a Gaussian random variable. We now consider the covariance operator associated with $Y_j, j \in [1, n]$:

$$\text{Cov}(Y_j, Y_k) = \text{Cov}(a_{ij}(W_i - \bar{W}_i), a_{pk}(W_p - \bar{W}_p)). \quad (16)$$

Invoking the bilinearity of the covariance operator, this reduces to:

$$\begin{aligned} \text{Cov}(Y_j, Y_k) &= a_{ij}a_{pk}\text{Cov}(W_i - \bar{W}_i, W_p - \bar{W}_p) \\ &= a_{ij}a_{pk}\text{Cov}(W_i, W_p) \\ &= a_{ij}a_{pk}\sigma^2\delta_i^p \\ &= \delta_j^k\sigma^2 \end{aligned} \quad (17)$$

where δ_j^k is the Kronecker delta, which is equal to 1 when $j = k$ and 0 otherwise, and the last passage is justified by the orthogonality of the matrix a_{ij} , which is characterized by orthonormal columns/rows. The new variables then constitute a set of i.i.d. Gaussian variables and can be treated as such in the Bayesian inversion process, notably in the definition of the likelihood function.

References

- Abbasi, M. and Gholami, A. (2017). Polynomial chaos expansion for nonlinear geophysical inverse problems. *Geophysics*, 82(4):R259–R268.
- Annan, A. P. (2005). GPR methods for hydrogeological studies. In *Hydrogeophysics*, pages 185–213. Springer.
- Arcone, S. A., Lawson, D. E., Delaney, A. J., Strasser, J. C., and Strasser, J. D. (1998). Ground-penetrating radar reflection profiling of groundwater and bedrock in an area of discontinuous permafrost. *Geophysics*, 63(5):1573–1584.
- Balanis, C. A. (2012). *Advanced Engineering Electromagnetics*. John Wiley & Sons.
- Blatman, G. and Sudret, B. (2011). Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of computational Physics*, 230(6):2345–2367.
- Bodin, T. and Sambridge, M. (2009). Seismic tomography with the reversible jump algorithm. *Geophysical Journal International*, 178(3):1411–1436.
- Cordua, K. S., Hansen, T. M., and Mosegaard, K. (2012). Monte Carlo full-waveform inversion of crosshole GPR data using multiple-point geostatistical a priori information. *Geophysics*, 77(2):H19–H31.
- Dietrich, C. R. and Newsam, G. N. (1997). Fast and exact simulation of stationary gaussian processes through circulant embedding of the covariance matrix. *SIAM Journal on Scientific Computing*, 18(4):1088–1107.
- Dou, Q., Wei, L., Magee, D. R., and Cohn, A. G. (2016). Real-time hyperbola recognition and fitting in GPR data. *IEEE Transactions on Geoscience and Remote Sensing*, 55(1):51–62.
- Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). Hybrid Monte Carlo. *Physics letters B*, 195(2):216–222.
- Ernst, J. R., Maurer, H., Green, A. G., and Holliger, K. (2007). Full-waveform inversion of crosshole radar data based on 2-D finite-difference time-domain solutions of Maxwell’s equations. *IEEE transactions on geoscience and remote sensing*, 45(9):2807–2828.
- Friedli, L., Linde, N., Ginsbourger, D., and Doucet, A. (2021). Lithological tomography with the correlated pseudo-marginal method. *Geophysical Journal International*, 228(2):839–856.
- Galetti, E., Curtis, A., Baptie, B., Jenkins, D., and Nicolson, H. (2017). Transdimensional Love-wave tomography of the British Isles and shear-velocity structure of the East Irish Sea Basin from ambient-noise interferometry. *Geophysical Journal International*, 208(1):36–58.
- Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical science*, 7(4):457–472.
- Giannakis, I., Giannopoulos, A., and Warren, C. (2019). A machine learning-based fast-forward solver for ground penetrating radar with application to full-waveform inversion. *IEEE Transactions on Geoscience and Remote Sensing*, 57(7):4417–4426.
- Giannakis, I., Giannopoulos, A., Warren, C., and Sofroniou, A. (2021). Fractal-constrained crosshole/borehole-to-surface full-waveform inversion for hydrogeological applications using ground-penetrating radar. *IEEE Transactions on Geoscience and Remote Sensing*.
- Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378.

- Good, I. J. (1992). Rational decisions. In *Breakthroughs in statistics*, pages 365–377. Springer.
- Goodman, J. and Weare, J. (2010). Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80.
- Grasmueck, M., Weger, R., and Horstmeyer, H. (2005). Full-resolution 3D GPR imaging. *Geophysics*, 70(1):K12–K19.
- Haario, H., Laine, M., Lehtinen, M., Saksman, E., and Tamminen, J. (2004). Markov chain Monte Carlo methods for high dimensional inversion in remote sensing. *Journal of the Royal Statistical Society: series B (statistical methodology)*, 66(3):591–607.
- Haario, H., Saksman, E., and Tamminen, J. (2001). An adaptive Metropolis algorithm. *Bernoulli*, pages 223–242.
- Hansen, T. M. and Cordua, K. S. (2017). Efficient monte carlo sampling of inverse problems using a neural network-based forward—applied to gpr crosshole traveltime inversion. *Geophysical Journal International*, 211(3):1524–1533.
- Hansen, T. M., Cordua, K. S., Jacobsen, B. H., and Mosegaard, K. (2014). Accounting for imperfect forward modeling in geophysical inverse problems—exemplified for crosshole tomography. *Geophysics*, 79(3):H1–H21.
- Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109.
- Higdon, D., McDonnell, J. D., Schunck, N., Sarich, J., and Wild, S. M. (2015). A Bayesian approach for parameter estimation and prediction using a computationally intensive model. *Journal of Physics G: Nuclear and Particle Physics*, 42(3):034009.
- Holm-Jensen, T. and Hansen, T. M. (2020). Linear waveform tomography inversion using machine learning algorithms. *Mathematical Geosciences*, 52(1):31–51.
- Hunziker, J., Laloy, E., and Linde, N. (2019). Bayesian full-waveform tomography with application to crosshole ground penetrating radar data. *Geophysical Journal International*, 218(2):913–931.
- Irving, J. and Knight, R. (2006). Numerical modeling of ground-penetrating radar in 2-D using MATLAB. *Computers & Geosciences*, 32(9):1247–1258.
- Khan, A., van Driel, M., Böse, M., Giardini, D., Ceylan, S., Yan, J., Clinton, J., Euchner, F., Lognonné, P., Murdoch, N., et al. (2016). Single-station and single-event marsquake location and inversion for structure using synthetic martian waveforms. *Physics of the Earth and Planetary Interiors*, 258:28–42.
- Kuroda, S., Takeuchi, M., and Kim, H. J. (2007). Full-waveform inversion algorithm for interpreting crosshole radar data: a theoretical approach. *Geosciences Journal*, 11(3):211–217.
- LaBrecque, D., Alumbaugh, D. L., Yang, X., Paprocki, L., and Brainard, J. (2002). Three-dimensional monitoring of vadose zone infiltration using electrical resistivity tomography and cross-borehole ground-penetrating radar. In *Methods in Geochemistry and Geophysics*, volume 35, pages 259–272. Elsevier.
- Laloy, E., Linde, N., Jacques, D., and Vrugt, J. A. (2015). Probabilistic inference of multi-Gaussian fields from indirect hydrological data using circulant embedding and dimensionality reduction. *Water Resources Research*, 51(6):4224–4243.

- Laloy, E., Rogiers, B., Vrugt, J. A., Mallants, D., and Jacques, D. (2013). Efficient posterior exploration of a high-dimensional groundwater model from two-stage Markov chain Monte Carlo simulation and polynomial chaos expansion. *Water Resources Research*, 49(5):2664–2682.
- Levy, S., Hunziker, J., Laloy, E., Irving, J., and Linde, N. (2021). Using deep generative neural networks to account for model errors in markov chain monte carlo inversion. *Geophysical Journal International*, 228(2):1098–1118.
- Li, J. and Marzouk, Y. M. (2014). Adaptive construction of surrogates for the Bayesian solution of inverse problems. *SIAM Journal on Scientific Computing*, 36(3):A1163–A1186.
- Lüthen, N., Marelli, S., and Sudret, B. (2021a). Automatic selection of basis-adaptive sparse polynomial chaos expansions for engineering applications. *International Journal for Uncertainty Quantification*.
- Lüthen, N., Marelli, S., and Sudret, B. (2021b). Sparse polynomial chaos expansions: Literature survey and benchmark. *SIAM/ASA Journal on Uncertainty Quantification*, 9(2):593–649.
- Madsen, R. B. and Hansen, T. M. (2018). Estimation and accounting for the modeling error in probabilistic linearized amplitude variation with offset inversion. *Geophysics*, 83(2):N15–N30.
- Marelli, S., Lüthen, N., and Sudret, B. (2021a). UQLab user manual – Polynomial chaos expansions. Technical report, Chair of Risk, Safety and Uncertainty Quantification, ETH Zurich, Switzerland. Report # UQLab-V1.4-104.
- Marelli, S. and Sudret, B. (2014). UQLab: A framework for uncertainty quantification in Matlab. In *Vulnerability, Uncertainty, and Risk (Proc. 2nd Int. Conf. on Vulnerability, Risk Analysis and Management (ICVRAM2014), Liverpool, United Kingdom)*, pages 2554–2563.
- Marelli, S., Wagner, P.-R., Lataniotis, C., and Sudret, B. (2021b). Stochastic spectral embedding. *International Journal for Uncertainty Quantification*, 11(2).
- Marzouk, Y. and Xiu, D. (2009). A stochastic collocation approach to Bayesian inference in inverse problems. *Communications In Computational Physics*, 6.
- Marzouk, Y. M., Najm, H. N., and Rahn, L. A. (2007). Stochastic spectral methods for efficient Bayesian solution of inverse problems. *Journal of Computational Physics*, 224(2):560–586.
- Meles, G., Greenhalgh, S., van der Kruk, J., Green, A., and Maurer, H. (2012). Taming the non-linearity problem in gpr full-waveform inversion for high contrast media. *Journal of Applied Geophysics*, 78:31–43.
- Meles, G. A., Van der Kruk, J., Greenhalgh, S. A., Ernst, J. R., Maurer, H., and Green, A. G. (2010). A new vector waveform inversion algorithm for simultaneous updating of conductivity and permittivity parameters from combination crosshole/borehole-to-surface GPR data. *IEEE Transactions on geoscience and remote sensing*, 48(9):3391–3407.
- Métivier, D., Vuffray, M., and Misra, S. (2020). Efficient polynomial chaos expansion for uncertainty quantification in power systems. *Electric Power Systems Research*, 189:106791.
- Nagel, J. B. (2019). Bayesian techniques for inverse uncertainty quantification. *IBK Bericht*, 504.
- Nielsen, L., Looms, M. C., Hansen, T. M., Cordua, K. S., Stemmerik, L., Miller, R., Bradford, J., and Holliger, K. (2010). Estimation of chalk heterogeneity from stochastic modeling conditioned by crosshole GPR traveltimes and log data. *Advances in near-surface seismology and ground-penetrating radar: SEG Geophysical Development Series*, 15:379–398.

- Olsson, O., Falk, L., Forslund, O., Lundmark, L., and Sandberg, E. (1992). Borehole radar applied to the characterization of hydraulically conductive fracture zones in crystalline rock. *Geophysical prospecting*, 40(2):109–142.
- Piscitelli, S., Rizzo, E., Cristallo, F., Lapenna, V., Crocco, L., Persico, R., and Soldovieri, F. (2007). GPR and microwave tomography for detecting shallow cavities in the historical area of “Sassi of Matera”(southern Italy). *Near Surface Geophysics*, 5(4):275–284.
- Rasmussen, C. E. (2003). Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer.
- Rawlinson, N. and Urvoy, M. (2006). Simultaneous inversion of active and passive source datasets for 3-D seismic structure with application to Tasmania. *Geophysical Research Letters*, 33(24).
- Sacks, J., Schiller, S. B., and Welch, W. J. (1989). Designs for computer experiments. *Technometrics*, 31(1):41–47.
- Slob, E., Sato, M., and Olhoeft, G. (2010). Surface and borehole ground-penetrating-radar developments. *Geophysics*, 75(5):75A103–75A120.
- Sochala, P., Gesret, A., and Le Maitre, O. (2021). Polynomial surrogates for bayesian traveltime tomography. *GEM-International Journal on Geomathematics*, 12(1):1–34.
- Stotzka, R., Wuerfel, J., Mueller, T. O., and Gemmeke, H. (2002). Medical imaging by ultrasound computer tomography. In *Medical Imaging 2002: Ultrasonic Imaging and Signal Processing*, volume 4687, pages 110–119. International Society for Optics and Photonics.
- Taillandier, C., Noble, M., Chauris, H., and Calandra, H. (2009). First-arrival traveltime tomography based on the adjoint-state method. *Geophysics*, 74(6):WCB1–WCB10.
- Tant, K. M. M., Galetti, E., Mulholland, A., Curtis, A., and Gachagan, A. (2018). A transdimensional Bayesian approach to ultrasonic travel-time tomography for non-destructive testing. *Inverse Problems*, 34(9):095002.
- Torre, E., Marelli, S., Embrechts, P., and Sudret, B. (2019). Data-driven polynomial chaos expansion for machine learning regression. *Journal of Computational Physics*, 388:601–623.
- Vu, M., Jardani, A., Revil, A., and Jessop, M. (2020). Magnetometric resistivity tomography using chaos polynomial expansion. *Geophysical Journal International*, 221(3):1469–1483.
- Wagner, P.-R., Fahrni, R., Klippel, M., Frangi, A., and Sudret, B. (2020). Bayesian calibration and sensitivity analysis of heat transfer models for fire insulation panels. *Engineering Structures*, 205:110063.
- Wagner, P.-R., Marelli, S., and Sudret, B. (2021a). Bayesian model inversion using stochastic spectral embedding. *Journal of Computational Physics*, 436:110141.
- Wagner, P.-R., Nagel, J., Marelli, S., and Sudret, B. (2021b). UQLab user manual – Bayesian inversion for model calibration and validation. Technical report, Chair of Risk, Safety and Uncertainty Quantification, ETH Zurich, Switzerland. Report UQLab-V1.4-113.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612.
- Warren, C., Giannopoulos, A., and Giannakis, I. (2016). gprmax: Open source software to simulate electromagnetic wave propagation for ground penetrating radar. *Computer Physics Communications*, 209:163–170.

- Warren, C., Giannopoulos, A., Gray, A., Giannakis, I., Patterson, A., Wetter, L., and Hamrah, A. (2019). A CUDA-based GPU engine for gprMax: Open source FDTD electromagnetic simulation software. *Computer Physics Communications*, 237:208–218.
- Xiu, D. and Karniadakis, G. E. (2002). The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM journal on scientific computing*, 24(2):619–644.
- Zhao, D., Lei, J., and Liu, L. (2008). Seismic tomography of the Moon. *Chinese Science Bulletin*, 53(24):3897–3907.